

**SafeTrip 21 Initiative:
Networked Traveler Foresighted Driving Field Experiment**

Draft Interim Report
December 27, 2010

Authors

**Christian Manasseh
Christopher Nowakowski
Daniel Vizzini
Raja Sengupta**

**Principal Investigator: Raja Sengupta
Lead Engineer: Christopher Nowakowski**

ABSTRACT

This report describes the SafeTrip-21, Networked Traveler Foresighted Driving Field Experiment. This experiment is conducted by UC Berkeley's California PATH in collaboration with researches from Science Applications International Corporation (SAIC) and Delcan Corporation, who served as an independent evaluation team during the course of the study. This report details the work and analysis that was performed by UC Berkeley. The report describes the test concept, the technology built, the experiment designed, and a preliminary analysis of the test data. The data is analyzed for correctness, integrity, and the potential of the technology to encourage safer driving.

1.0 INTRODUCTION

This report details one of the projects funded under the SafeTrip-21 initiative, the Networked Traveler Foresighted Driving Field Experiment. This experiment is conducted by UC Berkeley's California PATH in collaboration with researches from Science Applications International Corporation (SAIC) and Delcan Corporation, who served as an independent evaluation team during the course of the study. This report details the work and analysis that was performed by California PATH. A separate report has been produced by the SAIC/Delcan independent evaluation team (see Jasper, et al., 2010).

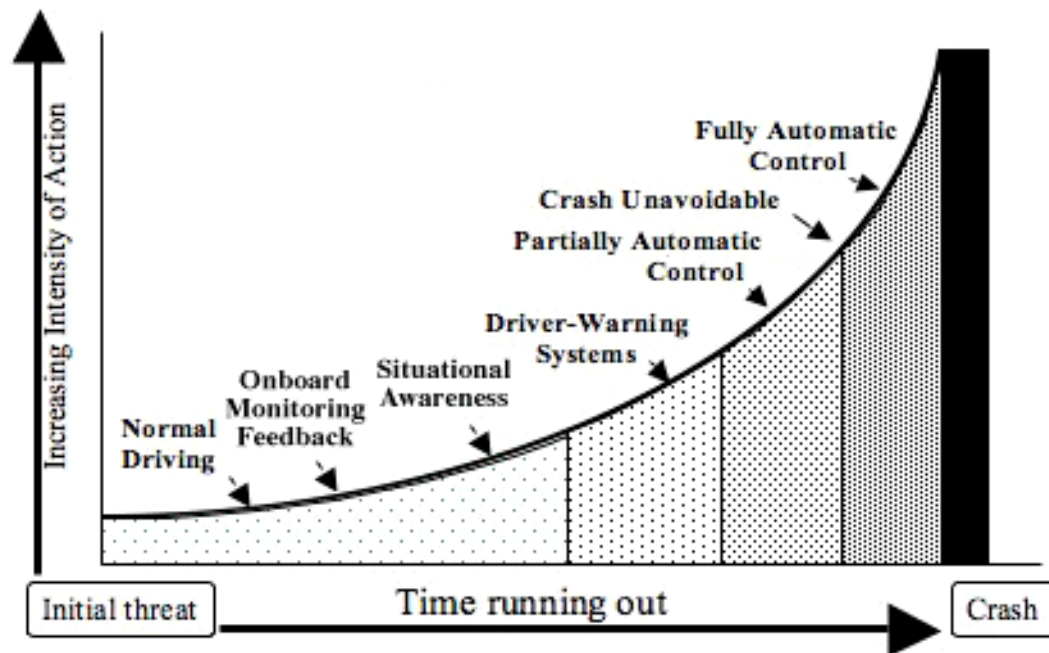
The goal of this project is focused on studying the effects of an Advanced Driver Assistance System (ADAS) providing soft-safety alerts regarding "Slow Traffic Ahead" when driving on a freeway. Currently, drivers may receive general traffic information about congestion through the internet, radio broadcasts, or the 511 hotline; however, the information given to drivers is fairly unspecific (such as travel times between two points or incidents along a route), and it is likely to be poorly timed. Drivers may only look at traffic information pre-trip, or they may get updates during their trip every 15 minutes or so, if they are tuned to a local radio station which provides such updates.

In evaluating the current systems for the dissemination of traffic information, it could be argued that there is a large gulf of execution between the needs of the user and the available sources of information. While pre-trip and periodic traffic information might support drivers' strategic decision making such as affecting their choice of route, this type of information is generally too coarse to affect tactical or operational decisions such as lane selection or choice of speed. However, this gulf of execution is not necessarily due to a lack of available data. Over the past decade or so, there has been a boom in number of organizations engaged in the collection, warehousing, and dissemination of real-time traffic information. In the San Francisco Bay Area alone, between the monitoring capabilities of two companies, Navteq and Speedinfo, nearly 5000 sensors exist on the freeway system that can provide traffic speed estimates with latencies that might be measured only in the terms of a few minutes.

Utilizing the available data that is already collected about the San Francisco Bay Area freeways, combined with the knowledge of an individual vehicle's position and speed, an ITS system was conceived that could provide drivers with real-time, individually tailored, traffic information. Early phases in the planning of this project explored how such a system might best be used to affect driving safety and reduce congestion, and this analysis determined that end-of-queue crashes were a significant problem which could be addressed by such a system. The remainder of the project focused on implementing a Bay Area wide prototype system and testing that system with naïve drivers.

2.0 The Foresighted Driving ADAS

The term ADAS (Advanced Driver Assistance System) is used to refer to any number of systems that might be conceived to support the driver in the task of driving such as backup cameras collision avoidance systems, blind spot warning systems, forward collision warning, etc. As first discussed in a NHTSA report (1992), crash countermeasures can be grouped or categorized along two axis, time and intensity. As shown in Figure 1, in any driving scenario which may lead to a crash, as time progresses, the intensity of the action required to prevent the crash increases. This figure has been adapted over time. The situational awareness category was added by Cody (2005), and the onboard monitoring feedback category was added by Misener, Nowakowski, Lu, et al. (2007).



Source: Figure adapted from National Highway Traffic Safety Administration, 1992.

Figure 1: Spectrum of ADAS Applications as a Function of Time and Intensity of Action.

Furthermore, over the years, there has been much debate as to where to implement partial or fully automatic vehicle control. As an example, in preventing rear-end collisions, some systems have been designed to mitigate crashes, only implementing fully automatic braking once a crash imminent situation has been detected. This is the case that is currently depicted in the graph, where fully automatic control is only implemented after a crash imminent scenario has been detected. The Lexus LS 430, introduced in the US in 2003, is one example of such a system. However, more recently, other systems have been designed to engage full braking with the intent to stop the vehicle before a crash occurs. The system introduced on the Volvo S80 and XC70 at the end of 2007 are examples of such a system. In the context of this project, this distinction is somewhat moot since we are only dealing with ADAS applications that fall further to the left on the graph.

Keeping with the example of preventing rear-end collisions (the most common type of end-of-queue crash), there has been much work published on the right side of this graph which represents what some might refer to as “hard safety” ADAS applications. In this context, the term “hard safety” is used to designate that immediate action, in the time frame of 1-3 seconds, is required by the driver to correct the situation. Hard safety applications include forward collision warning systems and road departure warning systems. In contrast, ADAS applications that fall more towards the left side of this graph might be referred to as “soft safety” applications. The term “soft safety” is used to designate that immediate action is not required by the driver. Soft safety applications are typically thought of as providing information that will be relevant to the driver in the range of 10-60 seconds from the point of delivery. Essentially, the scope of a soft safety ADAS application is fairly similar to the goals of most road signs – to provide the driver with preview about the road layout or hazards ahead.

Although there is certainly a continuum, and some applications may skirt the line between hard and soft safety, the differences between the two types of systems can also be summarized by their designed intent. Hard safety ADAS applications are typically designed to provide the driver with a warning that is intended to lead the driver towards a specific action that he or she is currently neglecting. In essence, hard safety systems typically act as a last-second intervention to prevent the driver from getting into a crash-imminent situation. Soft safety ADAS applications are typically designed to provide the driver with information that can be used to make better decisions regarding future events. Soft safety applications may provide the driver with enhanced situational awareness, or the case of the application that was prototyped and tested in this report, it may provide the driver with knowledge of traffic conditions ahead that may or may not be currently visible to the driver.

Concept

The experiment described in this report is centered around the testing of a foresighted driving ADAS concept that utilizing currently available traffic information and capable of providing targeted messages to individual drivers regarding traffic conditions that will soon be encountered. Thus, instead of just telling the driver that there may be congestion on some part of the freeway ahead, the foresighted driving system concept provides alerts to those drivers rapidly approaching an end-of-queue scenario.

The typical use case scenario developed for the foresighted driving ADAS concept starts off with a driver who travelling down a freeway at near free flow speeds of 55 to 65 mph. Unbeknownst to the driver, there is a queue of slowed traffic moving at 25 mph that is building 1 to 2 miles ahead, either due to a bottleneck in the freeway or an incident on the roadway. Furthermore, the end-of-queue may not be immediately visible to the driver due to curves or changes in grade. Without an ADAS system, a driver will continue traveling at their current speed until they notice the end-of-queue ahead, and then they will brake as it becomes necessary, following the trajectory of any cars in front of them and likely propagating any existing braking shock waves. Occasionally, some drivers might be surprised by the end-of-queue, which could lead to hard braking or, in the worst case, a secondary, end-of-queue crash.

Although the main end-of-queue crash problem is essentially a rear-end collision problem, the parameters of the use-case scenario laid out above differ slightly from the parameters used in most of the use-case scenarios for vehicle-based forward collision warning systems. Vehicle-based forward collision warning systems employ either a radar or lidar to detect the forward vehicle. These sensors have a limited range, which is a distinct disadvantage when the use case includes a very large speed differential, and like the driver, these sensors cannot see around blind curves. Additionally, many of the systems that rely on radar Doppler effects have a hard time seeing stopped vehicles, which are a distinct possibility in the end-of-queue crash scenario.

Thus, the conceptualized foresighted driving ADAS is able to detect the developing end-of-queue scenario at least 60-90 seconds before the driver reaches the end-of-queue, and provide the driver with an informational alert, preventing the potential for being surprised by the end-of-queue and reducing the number of hard braking vehicles near the end-of-queue. Theoretically, this could have the potential to reduce the number of secondary, end-of-queue crashes, and it could also provide minor benefits to traffic flow by smoothing traffic flow. However, there is currently little research to verify whether or not this type of soft safety alert would either alter a driver's behavior, or ultimately prove useful in either reducing end-of-queue collisions or helping to smooth traffic flow. This project was undertaken for these reasons.

The Systems Built

The Networked Traveler Foresighted Driving ADAS has a client-server architecture. The clients consist of 4 instrumented test vehicles, each containing their own onboard DAS. The server is a web server located at the California PATH Richmond Field Station. During the experiment, instrumented vehicles are driven by the participants, and these vehicles communicate with the server using cellular modems. Data is continually recorded both on the vehicles' DAS and on the server, so in effect, there are multiple sources of data being recorded in multiple locations during the experiment. Throughout the experiment, data from the various sources is periodically downloaded to a long-term data storage server for archival, processing, and analysis.

We equip 4 research vehicles (two 2008 Nissan Altimas and two 2007 Audi A3s) with an Advanced Driver Assistance System (ADAS) that provides drivers with an auditory, "Slow Traffic Ahead," alert. The alert sounds whenever the driver is rapidly approaching slowed or stopped traffic while driving on the freeway. The "Slow Traffic Ahead" ADAS is always on whenever the vehicle is running. From the driver's point of view the system functions as follows: When the driver is approaching a section of slowed traffic on the freeway, about 60 seconds before reaching that slowed traffic, the radio mutes, and the alert system plays a message that says something like "Slow Traffic Ahead, 35 mph," warning the driver that traffic ahead is moving significantly slower than he or she is currently travelling. Although the actual algorithm is a bit more complex, the alert generally triggers if the speed of traffic is 10-15 mph slower than the speed at which the subject is travelling.

The alert system works by transmitting a vehicle ID and the vehicle's location, speed, and direction over a wireless cellular 3G modem to a server located at California PATH. The server is continually tracking the state of traffic in the San Francisco Bay Area, and it transmits, back to the vehicle, a list of potential locations where the system may need to give the driver an alert. However, the ADAS on the vehicle decides when to actually give an

alert to the driver. The communication between the ADAS on the vehicle and the server at California PATH is secured through obscurity, meaning that one would not only need to find and intercept the communication on the internet, but one would also need to understand what the strings of numbers in the communication packets meant. This project does not use any standard protocols, and most of the values being transmitted by the ADAS on the vehicle have been scaled in obscure ways to create a compact message.

In addition to the “Slow Traffic Ahead” ADAS alert system, the vehicles each contain both a Data Acquisition System (DAS) and video recording systems onboard for the purpose of monitoring driver behavior and reactions to the alerts that are given. There are actually only two computers being installed in the vehicles. The first computer provides both the DAS and ADAS functions, and the second computer provides the video recording system functions. Audio is not recorded in this study.

The DAS records the following kinds of parameters:

- Vehicle throttle, braking, speed, acceleration, yaw rate, steering, and other vehicle parameters
- Vehicle GPS position
- Forward sensing of a lead vehicle, i.e., following distance and relative closing speed
- ADAS alert system parameters such as when and why an alert was given to the driver

The video recording system records includes the following cameras:

1. A camera viewing the forward road scene
2. A camera viewing the driver
3. A camera viewing the rear road scene
4. A video feed from the data recording system showing select parameters such as a timestamp, vehicle speed, and a visual indication of the onset of any auditory alerts that are given to the driver.

All of the data and video recording systems are nearly transparent to the driver. The two computers that make up the DAS, the ADAS, and the video recording are stored in the trunk of the vehicles. All cameras and other instrumentation are hidden from the driver’s view as best as possible. The only instrumentation readily visible to the driver is the camera focused on the driver’s face and a small speaker used to sound the foresighted driving alerts. Cameras are mounted to not interfere with the driver’s view, and cables are always hidden as best as possible. All of the systems are turn-key, meaning that all the driver needs to do is start the vehicle and start driving. The data and video recording systems and the “Slow Traffic Ahead” ADAS system automatically starts up and shuts down with the vehicle.

While a participant is in possession of one of the research vehicles, the data and video that are collected are stored in the DAS and video recording systems on the vehicle. The data collected on the vehicle is only downloaded from the onboard system when the vehicle is returned by the participant. The download process is manual, requiring a researcher to connect a USB hard drive to the DAS and initiate a download. The data and videos downloaded to the USB hard drive are then uploaded to a data storage and analysis server located in a secured room at California PATH’s Richmond Field Station facility. Any data on the USB hard drive is deleted after it has been successfully transferred.

The Foresighted Driving Data Server is primarily responsible for aggregating and preprocessing the traffic information feeds, purchased from SpeedInfo and Navteq, and then providing that information to the client vehicles upon request. In addition to this primary function, the server-side application records 3 types of data files during the experiment: raw traffic information, vehicle traces, and alert traces.

In addition to processing the traffic information feeds provided by SpeedInfo and Navteq, the server-side application also archives the raw data feeds for most of the experiment. The experiment began on 7/17/2010, but logging of the raw traffic data did not start until 7/26/2010. Logging of the raw traffic data continued through the end of the experiment on 11/15/2010.

The Foresighted Driving ADAS Server is developed to deliver two objectives. Its first objective is to act as a communication broker between the several sources of traffic data and the client application installed in the vehicle. Its secondary objective is to capture the driver traces and alerts and enable the driver to provide feedback in the form of web-based surveys while the experiment is underway.

In its role as a communication broker, the server receives live feeds of traffic information from two main sources: Navteq Links and SpeedInfo Sensors. The two sources expose a web service interface that the Foresighted Driving

ADAS server queries on a regular bases to update its information. This data is fused together to present a set of trigger points on the highways of the Bay area with known traffic speed measurements. Prioritization is given to SpeedInfo sensors over Navteq links in areas were both sources report traffic speed information. In areas were more than one sensor exists from the same source and report different measurements, the confidence of the measurement as reported by the source is used to select the traffic speed in that area. The Foresighted Driving ADAS server then presents a web service interface that can be consumed by the Client Application residing in the car. This service is a JSON interface providing the client application with a list of trigger points within 3km radius of the GPS coordinate submitted by the client. Each trigger point is represented with an ID, Speed Measurement at the sensor location, GPS Lon and GPS Lat, Heading and time of the reported measurement. To access the JSON service, the application is required to authenticate using a username and password.

In its role as a web site for the experiment, the Foresighted Driving ADAS server collects information from the client application in the car as each driver is driving. GPS traces with 1Hz measurements and alerts that are sounded to the driver get submitted to the server at a certain frequency (controlled by the client application) while the experiment is underway. A graphical representation of those traces and alerts is rendered by the server on a Google map. The user, provided with a username and password, can login to the website hosted on the server and see their previous traces and alerts. They are also given the opportunity to provide feedback on the alerts they received.

In addition to the above two roles, the server also acts as a data archival center for all the data of the experiment. All data received from the two sources of traffic speed data is archived throughout the period of the experiment, so is the data captured from the vehicle's client application – this excludes data collected by the vehicle DAS and not submitted to the server, which is later on stored separately on a different archiving system as mentioned above.

3.0 The Experiment

Each Foresighted Driving ADAS equipped car is handed out to a participant for a period of 2 weeks. The participants are asked to drive the car as they would drive their own vehicle normally, particularly on their daily commutes to and from work. After potential participants have been screened, qualifying participants are scheduled for a date to start their participation in the study. Participation has 4 major parts:

1. Driver orientation session (approximately 2 hours)
2. Baseline driving data (approximately 1 week)
3. Soft-safety alert enabled driving data (approximately 1 week)
4. Driver debriefing session (approximately 2 hours)

Additionally, as part of the protocol, when the vehicles are returned they are cleaned, inspected, data is downloaded, and any potentially identifying data left over from the previous participant, such as address book entries or recent destinations that were entered into the navigation system, is wiped clean.

The experiment sample is composed of 24 participants, 12 females and 12 males with ages ranging from 23 to 61 (mean 42, std dev 10.5). As per the study eligibility requirements, all of the participants have a clean driving record for at least 3 years, and none of the participants have a DUI on record. The table below details some of the participants' characteristics. The self-estimated annual average mileage of the participants ranged from 10,000 to 30,000 miles per year with a mean of 18,200 mi and a standard deviation of 5000 mi.

Table 3.: Test Participant Characteristics

#	Gender	Age	Est. Annual Mileage	Commute (miles)	Mean Commute Time (min)	Month of Participation (2010)
1	Male	38	20,000			July
2	Female	38	15,000			July
3	Male	43	15,000			July-August
4	Male	47	25,000			July-August
5	Female	35	15,000			August
6	Male	49	12,000			August
7	Female	52	18,000			August-September
8	Female	42	24,000			August-September
9	Male	61	25,000			September
10	Male	45	12,000			September
11	Female	37	10,000			September
12	Male	24	--			September
13	Female	47	15,000			September-October
14	Female	24	15,000			September-October
15	Female	23	25,000			September-October
16	Female	45	20,000			September-October
17	Male	50	15,000			October
18	Female	33	20,000			October
19	Male	52	15,000			October
20	Female	44	15,000			October
21	Male	44	18,000			October-November
22	Male	49	20,000			October-November
23	Female	27	20,000			October-November
24	Male	60	30,000			October-November

Most of the participants followed a two-week protocol as intended; however, the testing of two of the participants, one male and one female, deviated from the protocol slightly. In the case of the male participant, he originally started the protocol in August, but after the week of baseline driving, it was determined that there was a software error with the ADAS system on the vehicle which could not be fixed in the field. Since the participant had not seen any of the ADAS alerts, the vehicle was taken out of service and the participant was rescheduled to repeat both the baseline and testing weeks in a later session. In the case of the female participant, after the baseline week, it was determined that a configuration error had been made with the ADAS system. The error was corrected while the participant was still in possession of the test vehicle, and the baseline week was repeated, followed by the week of ADAS testing. In both cases, the participants simply ended up performing an additional week of baseline driving. We assume the deviations from the protocol do not affect the testing outcome.

4.0 Data Analysis To Date

4.1 Data Set Overview

Unlike a typical driving experiment, there is no set number of trials with the Networked Traveler Foresighted Driving ADAS. Instead, drivers are given free use of a test vehicle to drive on their own for a period of two weeks. Because the different vehicles are introduced into the experiment as they became ready, 16 of the participants drove Nissan Altimas and 8 of the participants drove Audi A3s. Furthermore, since only one Audi A3 had a manual transmission, 20 of the participants completed the testing with an automatic transmission, and 4 of the participants completed the testing with a manual transmission.

The concept of a trip, as discussed in this section, is defined as being from the time the vehicle ignition is started and the DAS started recording data (usually 1-2 minutes after the ignition is started) until the time that the ignition is switched off by the participant. It is expected that there will be at least 20 commuting trips, 10 during the baseline week and 10 during the week where the ADAS was enabled, but as shown in the table below the number of trips recorded by each participant ranged from a low of 35 to a high of 131 with a mean of 75 (± 25) trips per participant for a total of 1808 trips recorded by the four vehicles' DASs.

Table 4.1: Number of Trips Recorded by Participant.

Driver	Gender	Total Trips	Baseline Trips	ADAS Enabled Trips	No Driving Data Trips	Bad Data Trips
1	Male	102	42	47	13	0
2	Female	101	42	37	22	0
3	Male	99	29	35	35	0
4	Male	67	31	22	14	0
5	Female	38	16	21	1	0
6	Male	131	55	62	12	2
7	Female	62	26	26	9	1
8	Female	58	28	18	10	2
9	Male	78	23	41	13	1
10	Male	62	27	25	9	1
11	Female	60	23	31	4	2
12	Male	84	38	33	8	5
13	Female	70	31	28	11	0
14	Female	107	47	43	14	3
15	Female	75	35	25	15	0
16	Female	99	44	37	18	0
17	Male	114	42	47	25	0
18	Female	71	35	35	1	0
19	Male	84	38	46	0	0
20	Female	57	21	30	5	1
21	Male	43	25	18	0	0
22	Male	69	29	38	2	0
23	Female	42	5	19	1	17
24	Male	35	16	18	0	1
Total	--	1808	748	782	242	36

However, not all of the trips recorded are useful. As shown in the table, anywhere from 0 to 35 trips per participants, 242 trips in total, contain no useful driving data. This category consisted mostly of trips that were very short, and the only part of the trip recorded by the DAS is the vehicle maneuvering into a parking space or idling. Additionally, between 0 and 17 trips per participant, 36 trips in total, are categorized as having bad data. The ADAS system failures commonly occurred due to a lack of cellular/internet communications or due to errors in configuring the vehicle for each participant.

Overall, the experiment resulted in 748 trip during baseline testing weeks and 782 trips during the ADAS testing weeks with a mean of 31 (± 11) baseline trips and 32 (± 11) ADAS-enabled trips per participant. However, some participants have as few as 5 baseline or 18 ADAS-enabled trips, while other have as many as 55 baseline or 62 ADAS-enabled trips. As shown in the table below within those trips, a total of 1169 conditions are recorded during the baseline week where the Networked Traveler Foresighted Driving ADAS would have issued an alert, and between 715 and 732 times that the ADAS actually did issue an alert to the drivers.

Table 4.2 Foresighted Driving Alert Conditions by Participant.

Driver	Gender	Baseline Alerts (Muted)	Vehicle-Recorded ADAS Alerts	Server-Recorded ADAS Alerts
1	Male	70	72	70
2	Female	53	27	28
3	Male	77	38	40
4	Male	107	42	42
5	Female	19	18	17
6	Male	61	30	30
7	Female	53	14	16
8	Female	69	25	25
9	Male	32	64	66
10	Male	91	35	36
11	Female	26	27	28
12	Male	18	23	23
13	Female	60	41	44
14	Female	33	27	27
15	Female	78	38	38
16	Female	45	37	37
17	Male	18	14	14
18	Female	34	15	17
19	Male	23	24	26
20	Female	33	14	15
21	Male	29	11	11
22	Male	24	27	27
23	Female	55	22	25
24	Male	61	30	30
Total	--	1169	715	732

Currently, there are a number of discrepancies with the results provided in the above table that must be resolved before the analysis can continue. First, we would have expected the number of baseline alert-triggering conditions to roughly equal the number of actual alerts received by the drivers since the drivers were assumed to have similar driving patterns during the baseline and ADAS-enabled weeks. Instead, the number of baseline alert-triggering conditions was about 60% higher than the actual number of alerts received by the drivers. Although this needs to be investigated more thoroughly, it is likely that this difference will be accounted for by way that baseline alert-triggering conditions were reported. In effect, it appears likely situations where multiple baseline alert-triggering conditions were reported would have mapped to a single alert being issued. This is expected in places where navteq links are causing the alerts and the link lengths are small enough for a driver to traverse several links within 60 sec.

Second, in regards to the number of actual alerts received by the drivers, the Foresighted Driving ADAS alerts should have been recorded both on the vehicle DAS and by the server, since the ADAS reported all alerts back to the server through the cellular connection while the participant is driving. It was expected that the number of alerts recorded in each location should match, but it does not. At this point in time, further investigation is needed to explain the mismatched alert counts.

4.2 Geographic Data Examination

Our first classification of the data has been geographic. We conjectured this to be the best way to use our common sense to reveal any quality problem with the data. The first eight plots are representative examples from this exercise. These are chosen from over 200 such plots. Each plot is for a particular geographic location and summarizes the speed data for all drivers that drove through the location. The red lines represent data collected in the first week. During this week the DAS records where alerts would have sounded, but the system does not actually sound the alert. The table next to the figure calls these *non-audible alerts*. The blue lines represent the data collected in the second week, i.e., if the system determines an alert is to be sounded, the driver actually hears the alert. Thus the table accompanying the first figure indicates the red trace is an average of 19 traces and the blue line the average of 22 traces. Put another way, each speed time point on the red trace is the average of 19 speed values at the time. Likewise, each speed time point on the blue trace is the average of 22 speed values. The zero indicates the time at which the alert is sounded. The second row of the table records the number of drivers associated with the traces. Thus the 19 non-audible first week alert traces are produced by 6 different drivers and the 22 audible second week traces are produced by 8 drivers. The last column of this row states the total number of drivers. The value 10 is less than $8 + 6$ because some drivers are common to both the red and blue traces.


We are thus far able to only analyze the GPS data. Tools to examine the radar, video, or CAN data remain in development.

Location is defined as per the United States National Grid (USNG). The USNG is a coordinate system that divides the United States into quadrants of base-10 meters. It was modified slightly to specify 500-m by 500-m quadrants, and each trace was assigned a 500-m by 500-m grid based on their GPS location at the start of the alert. 500 m was deemed the largest possible size (and thus the largest possible sample size in a geographic unit) that is still small enough to isolate for geographic differences.

The eight plots included correspond to locations where we have data from 3 or more drivers. Most other locations are traveled by 2 subjects or less.


This geographic analysis helps us diagnose problems in the data before we analyze it for behavioral changes – the main objective of the experiment. For example, the first plot with 41 traces (the largest population of any quadrant) shows the alert sounded after the drivers slowed down. Clearly at this location, drivers had already reached the traffic queue before hearing the alert. To us this means, we failed to design the system correctly at this location. Thus this data set is excluded from the behavioral analysis discussed next. The 6th and 8th figures are also anomalous. In the 6th figure the drivers appear to be speeding up as they approach the alert time and reach a steady speed that is maintained even as the alert sounds. We are as yet unable to explain this behavior or the oscillations in the 8th plot. Thus these data sets have not yet been excluded from consideration. Once we are able to analyze the video data or radar data, we hope to get a better understanding of these two data sets. The preliminary behavioral analysis documented next is based on 1,395 traces.

The other five plots show what we expect when the system is designed correctly. In most cases, the alerts sounded before traffic slowed down. The other five graphs of the average speed of drivers before and after alerts sounded in the eight quadrants that had the greatest number of drivers that received both non-audible and audible alerts. The average longitude and latitude of GPS readings when the alerts began are included with an embedded Google Maps link.




Alert was rejected from subsequent analysis,
as traffic slowed before it sounded.
37.69861243, -121.9556681

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	19	22	41
Drivers in Alert	6	8	6




37.78810845, -122.2497619

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	13	8	21
Drivers in Alert	7	6	10




37.6769507, -122.118015

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	13	8	21
Drivers in Alert	7	5	9




37.88367166, -122.3086563

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	12	8	20
Drivers in Alert	5	6	8




37.86644883, -122.3038105

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	12	6	18
Drivers in Alert	4	4	5




Alert was not rejected from subsequent analysis,
although it appears to be a false alarm.
37.69088004, -122.0694342

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	7	8	15
Drivers in Alert	3	4	4



37.82534177, -122.2684474

	Non-Audible Alert (Red)	Audible Alert (Blue)	TOTAL
Alert Count	5	6	11
Drivers in Alert	3	5	5



Figures 4.1-4.8 Average speeds before and after an audible and non-audible alerts in the eight locations in which at least three drivers experienced audible and non-audible alerts

While the above charts are useful in qualitative analysis, most 500-m by 500-m quadrants lack sufficient sample sizes to perform any kind of rigorous statistical testing. Comparing responses for a given driver across all locations, instead of in a given location across all drivers, does provide adequate sample sizes, and is the method used below.

4.3 Preliminary Behavioral Analysis

Our first test statistic is speed. This is supported by the traffic calming literature. In order to better understand the challenges involved with this decision, we draw attention to the two traces of driver speed over time below:




Figure 4.9 Two traces of the speed of one driver two minutes before and after an alert is received. Note that the two traces are quite different, they both have similar values at time 0 and time 120, and similar maximum decelerations.

In Figure 4.9, the red, non-audible alert trace, shows a large variation in speed whereas the blue, audible alert trace, shows a gradual slowdown. Both traces approach the same final speed. This should be the case if traffic conditions are similar across the red and blue traces, i.e., in both cases drivers reach the speed of the queue.

The first test statistic used to capture the difference is the standard deviation of speed in the two minutes after an alert sounds:

$$SQRT(SUM(s-t - s\text{-bar})^2/2,400) \text{ for } t = 0 \text{ to } 2,400$$

where t is the index of 50-millisecond increments, with $t=0$ being the time the alert is sounded, s is speed of the vehicle recorded at this 50-millisecond increment. One would expect this standard deviation to be greater in the two minutes after a non-audible alert than after an audible alert.

A standard two sample t-test, with a null hypothesis that the standard deviations of speed in the two minutes after both an audible and non-audible alert were drawn from the same distribution of unknown mean and unknown variances was used to test the effects of the alert. This is a one-sided test, in which the alternate hypothesis is that the standard deviation in audible cases is lower than that of non-audible cases. The results are below:

<i>Driver</i>	<i>Alert Status</i>	<i>Test Statistic</i>		<i>N</i>	<i>F-Test</i>	<i>t-Test p-Value Assuming</i>	
		<i>Mean</i>	<i>SD</i>		<i>p-Value</i>	<i>Equal Variance</i>	<i>Unequal Variance</i>
18	Non-Audible	7.44	2.26	14	53.11%	0.00%	
18	Audible	2.39	1.89	14			
12	Non-Audible	7.05	2.31	12	72.61%	0.09%	
12	Audible	3.67	2.57	14			
7	Non-Audible	5.47	3.27	52	6.01%	0.26%	0.12%
7	Audible	3.49	2.34	29			
16	Non-Audible	5.13	3.14	36	25.25%	3.37%	
16	Audible	3.76	2.52	27			
19	Non-Audible	4.62	2.65	29	55.81%	3.93%	
19	Audible	3.17	2.27	15			
15	Non-Audible	6.03	3.85	22	18.89%	4.42%	
15	Audible	4.05	2.76	16			
8	Non-Audible	4.43	2.76	53	63.07%	5.06%	
8	Audible	3.08	2.42	14			
3	Non-Audible	5.47	2.94	77	92.37%	5.87%	
3	Audible	4.54	2.97	38			
1	Non-Audible	4.04	2.97	64	48.71%	6.81%	
1	Audible	3.30	2.72	69			
17	Non-Audible	5.94	2.75	43	69.84%	9.29%	
17	Audible	5.12	2.57	34			
14	Non-Audible	4.63	3.37	31	2.64%	15.77%	12.28%
14	Audible	3.72	1.97	17			
6	Non-Audible	5.84	3.27	12	73.64%	16.99%	
6	Audible	4.61	2.94	12			
13	Non-Audible	3.87	2.97	8	79.56%	27.48%	
13	Audible	3.12	2.80	16			
22	Non-Audible	2.55	3.14	21	48.10%	42.76%	
22	Audible	2.25	1.97	4			
2	Non-Audible	5.91	2.95	53	60.30%	43.61%	
2	Audible	5.79	2.66	24			
9	Non-Audible	4.10	2.62	40	32.35%	50.96%	
9	Audible	4.12	3.36	7			
10	Non-Audible	4.31	2.21	24	16.51%	59.16%	
10	Audible	4.47	2.87	57			
24	Non-Audible	5.42	1.78	23	19.90%	71.72%	
24	Audible	5.83	2.42	13			
25	Non-Audible	5.27	2.82	33	68.52%	71.76%	
25	Audible	5.78	3.05	16			
20	Non-Audible	3.69	3.86	13	87.12%	78.32%	
20	Audible	4.89	3.66	12			
21	Non-Audible	4.35	2.68	25	17.15%	78.46%	

21	Audible	5.28	3.82	9			
11	Non-Audible	4.00	3.14	73	49.21%	88.21%	
11	Audible	4.84	2.76	25			
23	Non-Audible	1.15	0.33	6	0.04%	96.01%	99.16%
23	Audible	3.01	2.40	13			
4	Non-Audible	3.15	2.75	98	73.72%	96.98%	
4	Audible	4.13	2.61	38			

Table 4.1 Results of the statistical testing on the standard deviation of speed in the 2-minutes after an audible and non-audible alerts. Test of unequal variance were only done for those driver's whose F-Test p-value was smaller than 10%.

The F-test is used to determine if the variances of the distribution for the non- audible alerts equals that for the audible alerts. In cases where its p-value is less than 10%, Excel's numerical approximation of the p-value of a t-test comparing two samples with unequal variance is used. In no cases is this p-value greatly different than that of the exact p-value of an equal-variance t-test.

While the low p-values of the F-tests warrant concern and further statistical testing, the results appear promising enough to merit more careful investigation. 10 of the 24 drivers have p-values under 10%, while only 2 have values greater than 90%. If these findings hold even after more rigorous analysis, it would suggest 10 of 24 drivers underwent a statistically significant behavior shift in a manner considered safer in the traffic calming literature.

REFERENCES

- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.
- Cody, D. (2005). Application of Situational Awareness to Driving: Design of Lane Change Assistant. *Proceedings of the 12th World Congress for Intelligent Transport Systems and Services*. San Francisco, California, November 6-10.
- Faul, F., Erdfelder, E., Lang, A.G., and Buchner, A. (2007). G*Power3: A flexible statistical power analysis program for social, behavioral, and biomedical sciences. *Behavior Research Methods* 39(2), 175-191.
- Jasper, K., Golembiewski, G., Armstrong, C., and Miller, S. (2010). *National Evaluation of the SafeTrip-21 Initiative: Final Report: Networked Traveler-Foresighted Driving* (Technical Report Contract Number: DTFH61-06-D-00005 Task Order T-09-008). Mclean, VA: Science Applications International Corporation (SAIC).
- Misener, J., Sengupta, R., Ahern, K., Datta Gupta, S., Dickey, S., Kuhn, T., Lian, T., Manasseh, C., Nelson, D., Rezai, S., Sarafsaleh, A., Shladover, S., and VanderWerf, J. (2010). *VII California: Development and Deployment Proof of Concept and Group-Enabled Mobility and Safety (GEMS)*, (Technical Report UCB-ITS-PRR-2010-26). Berkeley, CA: California PATH, Institute of Transportation Studies, University of California, Berkeley.
- Misener, J., Nowakowski, C., Lu, X., Koo, T., Margulici, J.D., Spring, J., Johnston, S., Kim, Z., Dickey, S., Kuhn, K., Kretz, P., Robin, J., and Walker, M. (2007). *Onboard Monitoring and Reporting for Commercial Motor Vehicle Safety Final Report* (Technical Report FMCSA-RRT-07-030). Washington, D.C.: Office of Research and Analysis, Federal Motor Carrier Safety Administration, U.S. Department of Transportation.
- National Highway Traffic Safety Administration. (2008). *Integrated Vehicle-Based Safety Systems (IVBSS) Phase I Interim Report* (Technical Report DOT-HS-810-952). Washington, D.C.: NHTSA, U.S. Department of Transportation.

National Highway Traffic Safety Administration. (2005). *Automotive Collision Avoidance System Field Operational Test Final Program Report* (Cooperative Agreement No. DTNH22-99-H-07019 Deliverable 42). Washington, D.C.: NHTSA, U.S. Department of Transportation.

National Highway Traffic Safety Administration. (1992). *NHTSA IVHS Plan* (Technical Report DOT-HS-807-850). Washington DC: NHTSA, Office of Crash Avoidance Research, U.S. Department of Transportation.